EXHIBIT 52

AMERICAN MARKETING ASSOCIATION

Article

Journal of Marketing Research
2021, Vol. 58(5) 888-907

DOI: 10.1177/00222437211030201
journals.sagepub.com/home/mrj
SAGE

First-Price Auctions in Online Display Advertising

Stylianos Despotakis, R. Ravi, and Amin Sayedi

Abstract

The authors link the rapid and dramatic move from second-price to first-price auction format in the display advertising market, on the one hand, to the move from the waterfalling mechanism employed by publishers for soliciting bids in a preordered cascade over exchanges to an alternate header bidding strategy that broadcasts the request for bid to all exchanges simultaneously, on the other. First, they argue that the move from waterfalling to header bidding was a revenue-improving move for publishers in the old regime when exchanges employed second-price auctions. Given the publisher move to header bidding, the authors show that exchanges move from second-price to first-price auctions to increase their expected clearing prices. Interestingly, when all exchanges move to first-price auctions, each exchange faces stronger competition from other exchanges, and some exchanges may end up with lower revenue than when all exchanges use second-price auctions; yet all exchanges move to first-price auctions in the unique equilibrium of the game. The authors show that the new regime hinders the exchanges' ability to differentiate in equilibrium. Furthermore, it allows the publishers to achieve the revenue of the optimal mechanism despite not having direct access to the advertisers.

Keywords

advertising, auctions, first-price auctions, game theory, header bidding, second-price auctions, waterfall bidding

Online supplement: https://doi.org/10.1177/00222437211030201

Advertising via online display ads has risen dramatically in the last decade. When display ads are generated by publishers, they are typically sold via ad exchanges to advertisers that bid in real time in spot auctions. In the last five years, most of these exchanges moved from a second-price to a first-price auction format. In this article, we show that this move can be attributed to another recent change in this industry regarding publishers' bid solicitation mechanism. Publishers moved from a sequential bid solicitation, known as waterfalling, to a parallel one, known as header bidding, in which they send the request for bid to all exchanges simultaneously. Using a parsimonious model of this ecosystem, we show that the publisher's move from waterfalling to header bidding can cause the subsequent move by the exchanges from second-price to first-price auctions. In addition to providing a new explanation for the change in auction format, our analysis suggests that this change reduces the ability of exchanges to differentiate themselves and lowers the fees they are able to extract from advertisers while allowing the publisher to achieve maximum revenue despite the advertisers being fragmented among multiple parallel exchanges.

Evolution of the Display Advertising Market

Display advertising, with an estimated market share of 54% in the United States, has grown to be a significant proportion of the digital advertising market.[1] The early promise of digital advertising came from search advertising that allowed advertisers to find customers at a deeper stage in their purchase funnel and also validate their interest by requiring payments only for clicks. However, as adoption of mobile devices grew and video content became more popular, a large volume of user attention became available in the form of user visits to websites

[1] See https://forecasts-na1.emarketer.com/584b26021403070290f93a56/5851918a0626310a2c1869ca.

Stylianos Despotakis is Assistant Professor of Marketing, City University of Hong Kong, Hong Kong (email: sdespota@cityu.edu.hk). R. Ravi is Andris Zoltners Professor of Business and Professor of Operations Research and Computer Science, Tepper School of Business, Carnegie Mellon University, USA (email: ravi@cmu.edu). Amin Sayedi is Associate Professor of Marketing, Foster School of Business, University of Washington, USA (email: aminsa@uw.edu).

and mobile apps other than search. This has led to an increase in the availability of user eyeballs, referred to as "impressions," on websites and mobile apps visited by users. Furthermore, new technology in display advertising, such as real-time bidding (RTB), has allowed advertisers to target consumers dynamically and at an individual level and has made display advertising more appealing than before.

Early methods of selling display ads involved more traditional channels with a sales force and the sale of fixed large inventories of eyeballs, and these were the mainstay of companies that ran publishing networks that produced such large streams of impressions (e.g., Yahoo, Microsoft). After traditional sales, the uncertain inventory of impressions that were unsold, the so-called remnants, were then auctioned off in real time in one or multiple marketplaces, called "exchanges." In RTB auctions, when such impressions became available, the publisher generated a request for bids dynamically and tried to sell the impression to the highest bidder. Given the lower volume of these available remnants, early publisher networks that processed remnant inventory preferred to send the request for bids to only a few reliable large advertising networks or exchanges so as to get a quick and reasonably high bid for the impression. Typically, such preferred exchanges were ordered in sequence of their expected price they fetched per impression, and the RTB system would sequentially go down this order of exchanges, generating a request and then waiting for a short while before timing out and moving to the next exchange in the sequence. The first acceptable bid in the order was accepted, and auction reserve prices (often called "floor prices" in this context) were used to control the level of acceptability. This form of RTB that evolved in the early days of display advertising was termed "waterfalling" (Zawadziński 2018).

One of the undesirable features of the waterfall model is that it creates a fractured market among the ad exchanges, leaving advertisers in a quandary about which exchanges to associate with to spend their budgets most effectively. In particular, the advertiser's decision of which exchange to join should take into account both its order in the waterfall and the competition within that exchange. For publishers, this format can lead to the loss of high-valued advertisers in the later stages of the waterfall, whose bids are never considered.

Given these inefficiencies, around 2014, publishers introduced a new format for requesting bids from ad exchanges, called "header bidding." In this format, rather than go through different partner ad exchanges in sequence, the publisher broadcasts the request for bid simultaneously to all ad exchanges and, after collecting all returned bids within a reasonable time frame, picks the best one.[2] Since its introduction, header bidding caught on very rapidly and became the mainstream format of publishers by the end of 2016. By some estimates, the percentage of the top publishers that used header bidding increased from 0% to over 70% in the period 2014–2016 (Insider Intelligence 2016).

[2] For this reason, this was also called "advance bidding" or "prebidding."

Before header bidding was introduced in the display advertising marketplace, the auction format for selling display ads was the well-established second-price format (with a potential reserve price set by the publisher), which the industry inherited from the paid search advertising world. However, in early 2017, right after the introduction of header bidding, several ad exchanges began experimenting with a first-price auction format instead. This move came about in a variety of ways, including the introduction of "soft floors" that were set by the ad exchanges. While the publisher supplied a reserve price with the request for bids, called the "hard floor," each ad exchange would set another, higher value as a soft floor and change the rule of the local auction in the following way: if there were at least two bids above the soft floor, they participated in a regular second-price auction; with only one bid above the soft floor, the soft floor then served as the clearing price; with all bids below the soft floor but some still above the hard floor, the bids participated in a first-price auction. Note that by setting the soft floor sufficiently high, the auction format is effectively converted from a second-price to a first-price auction. Indeed, several exchanges such as AppNexus advised advertisers to bid in soft-floor auctions just as they bid in first-price auctions (Gubbins 2017). The lack of transparency about the values of the soft floors set in these auctions led to such intermediate formats being quickly replaced by the more transparent first-price format with a reserve price (Shterev 2020). After Google's move to first-price auctions in 2019, all major exchanges now use first-price auctions to sell display advertising impressions, when a publisher sends the request for bid to multiple exchanges.

We emphasize that the move from second-price to first-price auctions happened in situations when a publisher sends the request for bid to multiple third-party exchanges (typically through header bidding) and allocates the impression to the exchange with the highest clearing price—for example, when exchanges such as Rubicon (Kershaw 2017), PubMatic (2021), and OpenX (2019) compete with each other to sell the publisher's impression. This is in contrast with situations in which the publisher sells its inventory directly to advertisers without calling third-party exchanges (e.g., when Google sells YouTube impressions) or uses only a single exchange to sell the impressions. Indeed, as observed in Tunuguntla and Hoban (2021), second-price auctions remain the standard when the publisher sells the impressions exclusively through one exchange, and, as such, the exchange is the final arbiter of impression placement (e.g., scenarios when OpenX runs the final auction [OpenX 2019]).

Research Questions

Our primary research question pertains to the rapid transition of exchanges from second- to first-price auction format, and why this move occurred at this late stage rather than with the initial advent of display advertising.

A second research question is the consequence of this move from second- to first-price auctions for the strategies of the publishers and the advertisers. Publishers will have to reevaluate

their floor prices as a consequence of this double move from waterfalling to header bidding and from second-price to first-price auctions. Similarly, advertisers will have to modify their bidding strategies from the old regime to the new. In addition, given the change in the marketplace rules, they may also need to reevaluate whether the choice of ad exchanges with which they affiliate themselves (often at nonnegligible costs) is optimal and worth the corresponding fees.

A third research question involves how these two changes in the marketplace (from waterfall to header bidding, and of the auction format from second to first price) affect ad exchange revenues, in both the short and the long term.

Contributions

In this article, we study these research questions using a simple model of the display advertising ecosystem. Our model involves a single publisher, two ad exchanges, and a minimal set of (four) advertisers that decide to affiliate with one of the ad exchanges by paying their respective fees. We begin our analysis in the old regime, where the ad exchanges use second-price auctions with reserve prices, and the publisher uses waterfalling. In this waterfalling setting, we show (in Proposition 1) that the revenue of the ad exchanges and the advertisers' utilities are not affected by the auction format, so there is no incentive for the exchanges to change from their historically prevalent second-price format. We note in Proposition 2 that in this setting, the ad exchanges are able to set nonzero entry fees for the advertisers.

Our next result (Proposition 3) shows that when the ad exchanges use a second-price auction, moving from waterfalling to header bidding increases the publisher revenue, thus providing a simple economic explanation for this initial move of the publishers.

To answer the primary research question of the subsequent move of the ad exchanges from second to first-price auctions, we analyze the choice of auction format by the exchanges after the publisher's move to header bidding. Our main result is a new explanation in Proposition 4, where we show that under header bidding, there is a unique equilibrium in which both exchanges run first-price auctions. We can contrast this to our previous observation (Proposition 1) that under waterfalling, the revenues of the ad exchanges are not affected by the auction format. In this way, we show that the move to first-price auctions might be a direct economic consequence of the widespread adoption of header bidding by publishers.

This result provides an alternate reason to the main explanation that has been advanced so far for this change in auction format in the literature (e.g., Akbarpour and Li 2020), which argues that the move is the result of trust issues because advertisers do not have to trust the exchange in a first-price auction. Our model gives an intuitive explanation for the rapid adoption of first-price auction by exchanges that occurred soon after the exponential growth of header bidding (Sluis 2019). As additional evidence of the plausibility of our explanation, we note that Google adopted the first-price auction format for its exchange platform, where it acts as an intermediary and the participating publishers use header bidding; however, it has retained the second-price auction format for the sales of its own inventory (e.g., YouTube), where it assumes the role of the publisher and there are no intermediaries (Cox 2019).

Next, we show that, under header bidding with first-price auctions, the exchanges' equilibrium fees for the advertisers become zero (Proposition 5). In other words, under waterfalling, the exchanges can differentiate themselves on the basis of their positions in the waterfall sequence, and under second-price auctions, the exchanges can differentiate based on the set of advertisers with which they are affiliated. This led to nonzero fees for ad exchanges under waterfalling with second-price auctions (Proposition 2). However, the combination of header bidding and first-price auctions removes the exchanges' ability to differentiate on the basis of their set of advertisers or their position and thus lowers their equilibrium buyer-side fees. This finding is consistent with recent reductions in exchange fees from an average of 25% in 2016 to around 15% in 2018, and these are predicted to be in single digits in near future (Sluis 2017a, c).

From a managerial point of view, our results shed light on how the new selling mechanism (i.e., the combination of header bidding and first-price auctions) affects the strategies of advertisers, publishers, and exchanges. We show that while advertisers should shade their bids in first-price auctions, they should bid as if all advertisers (from all exchanges) are in the same auction. In other words, under header bidding with first-price auctions, each advertiser is directly competing with all advertisers from all exchanges. For publishers, we show that the new mechanism greatly simplifies the reserve price optimization problem. Furthermore, by setting the reserve prices optimally, publishers can achieve the revenue of the optimal mechanism (Myerson 1981), even though they do not have direct access to advertisers. Finally, because the new mechanism eliminates the exchanges' ability to differentiate on the basis of the number of their advertisers and their position in the waterfall, exchanges have to devise new differentiation strategies to survive in the long run.

The rest of the article is structured as follows. First, we review the related literature and present the model. Then, we analyze the model, discuss the results, and conclude. All proofs are relegated to Appendix A.

Related Literature

Our work is related to the growing literature on online advertising auctions. Katona and Sarvary (2010) and Jerath, Ke, and Long (2011) study advertisers' incentives in obtaining lower versus higher positions in search advertising auctions. Sayedi, Jerath, and Srinivasan (2014) investigate advertisers' poaching behavior on trademarked keywords and their budget allocation across traditional media and search advertising. Desai, Shin, and Staelin (2014) analyze the competition between brand owners and their competitors on brand keywords. Lu, Zhu, and Dukes (2015) and Shin (2015) study budget constraints

and budget allocation across keywords. Zia and Rao (2019) look at the budget allocation problem across search engines. Wilbur and Zhu (2009) find the conditions under which it is in a search engine's interest to allow some click fraud. Cao and Ke (2019) and Jerath, Ke, and Long (2018) study manufacturer and retailers' cooperation in search advertising and show how it affects intra- and interbrand competition. Amaldoss, Desai, and Shin (2015) show how a search engine can increase its profits and also improve advertisers' welfare by providing first-page bid estimates. Berman and Katona (2013) study the impact of search engine optimization, and Amaldoss, Jerath, and Sayedi (2015) analyze the effect of keyword management costs on advertisers' strategies. Katona and Zhu (2017) show how quality scores can incentivize advertisers to invest in their landing pages and to improve their conversion rates. Long, Jerath, and Sarvary (2018) study the informational role of search advertising on the organic rankings of an online retail platform. Our work differs from these articles because we study display advertising auctions in RTB. In our article, the auctioneer (i.e., the exchange mechanism) is different from the publisher, whereas in search advertising models, the publisher (i.e., the search engine) also designs the auction.

Our work contributes to the vast literature on display advertising. Empirical works in this area have assessed the effectiveness of display advertising in various contexts. Lambrecht and Tucker (2013) demonstrate that retargeting may not be effective when consumers have not adequately refined their product preferences. Hoban and Bucklin (2015) find that display advertising increases website visitations for a large segment of consumers along the purchase funnel, but not for those who had visited before. Bruce, Murthi, and Rao (2017) examine the dynamic effects of display advertising and show that animated (vs. static) ads with price information are the most effective in terms of consumer engagement. Rafieian and Yoganarasimhan (2021) study the role of targeting in online advertising and show that ad networks may benefit from preserving customer privacy. Rafieian (2019) shows that publishers can improve their revenue by optimally sequencing the ads that they show to a customer in a session. On the theoretical front, Sayedi, Jerath, and Baghaie (2018) study advertisers' bidding strategies when publishers allow advertisers to bid for exclusive placement on the website. Zhu and Wilbur (2011) and Hu, Shin, and Tang (2015) study the trade-offs involved in choosing between "cost-per-click" and "cost-per-action" contracts. Berman (2018) explores the effects of advertisers' attribution models on their bidding behavior and their profits. Despotakis, Ravi, and Srinivasan (2021) and Gritckevich, Katona, and Sarvary (2018) look at how ad blockers affect the online advertising ecosystem, and Dukes, Liu, and Shuai (2019) show how skippable ads affect publishers' and advertisers' strategies as well as their profits. Kuksov, Prasad, and Zia (2017) study firms' incentives in hosting the display ads of their competitors on their websites. Choi and Sayedi (2019) study the optimal selling mechanism when a publisher does not know, but benefits from learning, the performance of advertisers' ads. These works, unlike ours, do not study the roles of intermediaries (i.e., exchange platforms) in the market. In contrast, the focus of our research is to study what triggered the intermediaries' move from second-price to first-price auctions and how this move affects publishers and advertisers in this market.

In the context of RTB auctions, Johnson (2013) estimates the financial impact of privacy policies on publishers' revenue and advertisers' surplus. Rafieian (2020) characterizes the optimal mechanism when the publisher uses dynamic ad sequencing. Zeithammer (2019) shows that introducing a soft reserve price, a bid level below which a winning bidder pays their own bid instead of the second-highest bid, cannot increase publishers' revenue in RTB auctions when advertisers are symmetric; however, it can increase the revenue when advertisers are asymmetric. The model in Zeithammer (2019) has only one exchange and, therefore, cannot distinguish between header bidding and waterfalling. We show that while under waterfalling the results of Zeithammer (2019) continue to hold, under header bidding first-price auctions generate a higher revenue for the publisher than second-price auctions, even with symmetric advertisers. Sayedi (2018) analyzes the interaction between selling impressions through RTB and selling through reservation contracts and shows that, to optimize their revenue, publishers should use a combination of RTB and reservation contracts. In Sayedi, there is only one exchange, and header bidding (compared with waterfalling) affects how advertisers in RTB compete with those in reservation contracts. Choi and Mela (2018) study the problem of optimal reserve prices in the context of RTB and, using a series of experiments, estimate the demand curve of advertisers as a function of the reserve price. Because the data set in Choi and Mela (2018) is from 2016, exchanges still use second-price auctions. We show that when exchanges use first-price auctions, the publisher's problem of reserve price optimization becomes much simpler. Choi et al. (2020) provide an excellent summary of the literature and key trends in the area of display advertising markets. They also mention the move of ad exchanges from second-price to first-price auctions to be related to header bidding in that it enables the highest bidder to win (which is not necessarily the case with the second-price format), but they do not provide a model or analysis. Interestingly, they leave it to future research to analyze the impact of the recent changes in selling mechanisms on advertisers' and publishers' revenues, which is a gap we attempt to fill with our work.

Model

There is one publisher that is selling an impression, two exchanges, and four advertisers that can bid for the impression through one of the exchanges. Each advertiser's valuation for the impression is an i.i.d. draw from a uniform distribution on the interval [0, 1], with cumulative distribution function $F(x) = x$.

Exchanges

Exchanges are intermediaries that connect publishers to advertisers. The revenue of an exchange comes from buyer-side fees

(i.e., how much they charge advertisers for their service) and seller-side fees (i.e., how much they charge publishers for their service). We assume that Exchange i, $i \in \{1, 2\}$, sets a fee $f_i \geq 0$ for advertisers that want to use its platform. In practice, the buyer-side fee can have a complex structure and be a combination of advertisers' bidding and winning volumes, as well as their average submitted bids.[3] In the interest of parsimony, we assume that an advertiser has to pay a flat fee f_i if it wants to use Exchange i, where f_i is set by the exchange. Given the fees and other parameters (discussed subsequently), the advertisers decide which exchange to join. We use n_i to denote the number of advertisers that use Exchange i. The exchanges also charge publishers a seller-side fee. Seller-side fees are negotiated between exchanges and publishers and are typically a fraction of an exchange's contribution to the publisher's revenue. In our model, we assume that the publisher pays fraction f of the revenue that it collects through Exchange i to Exchange i. For example, if an exchange sells the impression of the publisher for a price of 1, the publisher keeps $1 - f$ and gives f to the exchange. To facilitate exposition, we assume that seller-side fee f is the same for both exchanges and exogenous in the model.[4] If Exchange i sells the impression at price p, its total revenue is $n_i f_i + fp$, and if it does not sell the impression, its total revenue is $n_i f_i$.

In addition to setting buyer-side fees, the exchanges also decide what auction format to use. An exchange can use a first-price auction or a second-price auction to sell the publisher's impression on its platform. The exchange uses the format that maximizes its revenue; the format of the auction is revealed to the advertisers before they submit their bids. In both formats, the highest bidder wins as long as the bid is greater than or equal to the reserve price. The clearing price of a second-price auction is the maximum of the second-highest bid and the reserve price. The clearing price of a first-price auction is the highest bid. If no one bids the reserve price or higher, in both auction types, the clearing price is zero.

Publisher

When an impression arrives (i.e., a consumer visits the publisher's website or app), the publisher sends a "request for bid" to the exchanges. The publisher can send the request for bids to both exchanges simultaneously or sequentially. As we have discussed, the sequential strategy is called "waterfalling" and the simultaneous strategy is called "header bidding." Under waterfalling, the publisher waits for the outcome of the first exchange, and if the impression is sold in the first exchange (i.e., there is at least one bid greater than or equal to the reserve price), the publisher does not send it to the second exchange. If the impression is left unsold in the first exchange, the publisher sends it to the second exchange. The publisher can also choose the order of the exchanges (i.e., to which exchange to send the impression first); without loss of generality, we assume that, if the publisher uses waterfalling, it sends the impression to Exchange 1 first. Under header bidding, the publisher sends the impression to both exchanges at the same time. Each exchange runs an auction and sends its clearing price back to the publisher; the publisher selects the exchange with the highest clearing price as long as at least one of the clearing prices is greater than zero. If both clearing prices are zero (i.e., the impression is left unsold in both exchanges), the impression remains unallocated, and the publisher's revenue becomes zero.

When RTB started, waterfalling was the only strategy available to the publishers. In 2014, some publishers moved to header bidding, and by the end of 2016, more than 70% of top publishers in the United States were using header bidding. Even though the choice of header bidding versus waterfalling is a publisher's decision, in our model, we analyze the two models separately. This enables us to highlight how the publisher's move from waterfalling to header bidding triggered the adoption of first-price auctions by exchanges and explain how this market has evolved over time. Finally, the publisher sets reserve prices for each exchange. We use r_i to denote the reserve price of Exchange i. In practice, and also in our model, optimizing the reserve prices is an essential part of revenue optimization for publishers in RTB markets (Choi and Mela, 2018). If the impression is allocated to Exchange i, with clearing price p, the publisher's revenue is $(1 - f)p$.

Advertisers

There are $n = 4$ advertisers in our model. While we can prove many of our results for larger number of advertisers, having $n = 4$ has two benefits for us. First, the number is large enough so that if two advertisers join each exchange, we still have within-exchange competition between advertisers in both exchanges. Furthermore, we can prove the results analytically in the case of $n = 4$, whereas finding advertisers' bidding strategies for larger values of n becomes analytically intractable.[5] Advertisers make two decisions in our model. First, they decide which exchange to join, if any.[6] This decision happens after the advertisers learn the buyer-side fees, f_1 and f_2, but before they learn their valuation for the impression. In

[3] It has even been reported that many exchanges were not transparent in terms of what fees they charged advertisers, and in many cases advertisers were surprised when they realized that some of payments were being issued to the exchange instead of the publisher. For example, see Sluis (2017b).

[4] Note that even though the exchanges are ex ante symmetric, they can create value (and differentiate) by offering additional bids from new advertisers. In other words, the publisher is not constrained to work with only one exchange and benefits from allowing as many exchanges as possible, as long as each exchange brings new advertisers to the game. This allows the exchanges to set positive seller-side fees.

[5] In the Web Appendix, we numerically verify our results for larger numbers of advertisers.

[6] While it is possible for an advertiser to bid in several exchanges at the same time, this is not commonly observed in practice. Some potential reasons may be to avoid paying access fees to multiple platforms and to avoid the potential risks of indirectly competing against itself from different exchanges.

practice, advertisers choose an exchange before impressions arrive; the partnership between exchanges and advertisers is usually long term, and advertisers cannot switch exchanges in real time, before or after every impression. Therefore, advertisers have to use the valuation distribution, as opposed to the actual realization of the valuation, when deciding which exchange to join. An advertiser can also decide to not join either of the two exchanges (e.g., if the fees are too high). In that case, the advertiser's utility becomes zero (i.e., the advertiser "leaves the game").

When an impression arrives at an exchange, advertisers in that exchange decide how much to bid for the impression. At this time, each advertiser knows its private value for the impression, an i.i.d. draw from the uniform U[0, 1] distribution. The advertiser also knows the reserve prices, r_1 and r_2, and the format of the auction in both exchanges (i.e., whether each exchange uses first-price or second-price auctions). The advertiser, however, does not know other advertisers' valuation for the impression (but only their distribution). These assumptions are consistent with what advertisers know when submitting their bids in this market. Note that an advertiser's bid is only submitted to the exchange that it has joined. Under header bidding, winning in the affiliated exchange does not imply that the advertiser will get the impression, as the impression may be allocated to the winner of the other exchange. If an advertiser with valuation v, which uses Exchange i, is allocated an impression at clearing price p, the advertiser's utility is $v - p - f_i$. If the advertiser does not win, its utility is $-f_i$.

Timeline

Following is a summary of the timeline of the game. To highlight the similarities and for succinctness, we present both the waterfalling and header-bidding scenarios of the publisher in the same outline.

1. Exchanges decide their buyer-side fees f_1 and f_2.
2. Advertisers choose which exchange to join (if any).
3. The publisher sets reserve prices r_1 and r_2 for the two exchanges.
4. Exchanges decide their auction format (i.e., whether to use a second-price or first-price auction).
5. Advertisers' valuations are privately realized; they submit their bids to their affiliated exchanges.
6. Under waterfalling:
    a. Exchange 1 runs its auction. If the clearing price p_1 of the auction is larger than zero (i.e., at least one bidder bids at least the reserve price), then the publisher allocates the impression to the winner of Exchange 1 for price p_1 and the game ends.
    b. If the clearing price in Exchange 1 is zero (i.e., no bidder in Exchange 1 bids at least the reserve price), then the publisher moves to Exchange 2. The second exchange runs its auction with reserve price r_2 and sends the clearing price p_2 to the publisher. If $p_2 > 0$ (i.e., there is at least one bid greater than or equal to r_2), the publisher allocates the impression to the winner of Exchange 2 for price p_2. Otherwise, the impression remains unsold.
7. Under header bidding:
    a. The publisher sends the impression to both exchanges; the exchanges run their auctions simultaneously. They send their clearing prices p_1 and p_2 to the publisher.
    b. If $\max(p_1, p_2) > 0$ (i.e., in at least one exchange one bidder bids greater than or equal to the reserve price of that exchange), the publisher allocates the impression to the exchange with the higher clearing price[7] at that exchange's clearing price. Otherwise, the impression is left unsold.

Note that, in our timeline, the exchanges decide the auction format after the advertisers join exchanges and the publisher sets the reserve prices. This is because, in practice, during the transition from second-price to first-price auctions, some exchanges (such as Rubicon) announced that they would decide the auction format in real time, after an impression arrives (Gubbins 2017). While all exchanges eventually moved to a pure first-price auction format (as it happens in our model as well), from a modeling perspective, we have taken into account that they had the option of choosing a different format for every impression.

Rather than model the industry transition, an alternate motivation may be to analyze the current situation in the market where the exchanges commit to an auction format and announce it before the reserve prices are set. Motivated by this, and to verify the robustness of our results, in Appendix A we consider two alternative timelines in which exchanges decide the auction formats before the publisher sets the reserve prices. We first consider the case when advertisers join exchanges before the exchanges decide their auction format and then consider the case when advertisers join exchanges after the exchanges decide their auction format and before the publisher sets the reserve prices. We show that our main results continue to hold under both these extensions.

Analysis

We use backward induction to solve the game under each scenario, waterfalling and header bidding, separately. To keep the flow of this section consistent with the evolution of the display advertising industry, we start by analyzing the waterfalling game.

Waterfalling

In the last stage of the game, advertisers have to decide how much to bid for the impression. Suppose that there are n_1 and

[7] In case of a tie between the two exchanges, the winner is selected randomly with probability 1/2.

n_2 advertisers in the first and the second exchange and the reserve prices are set to r_1 and r_2, respectively.[8] The following lemma summarizes the advertisers' bidding strategies.

Lemma 1: Under waterfalling, if Exchange i is using a second-price auction with reserve r_i, all advertisers bid truthfully.[9] If the exchange is using a first-price auction with reserve r_i, an advertiser with valuation $v \geq r_i$ bids

$$v - \frac{v}{n_i} + \frac{r_i^{n_i}}{n_i \times v^{n_i - 1}}.$$

Lemma 1 summarizes the advertisers' bids in first-price and second-price auctions as a function of the reserve price and the number of advertisers. The expressions that we have for the bids are the standard expressions for first-price and second-price auctions with reserve prices (e.g., see Krishna 2009). The lemma shows that, under waterfalling, existence of Exchange 1 does not directly affect the bids of Exchange 2, and vice versa. In other words, advertisers in different exchanges do not directly compete with each other under waterfalling. These advertisers, however, compete with each other indirectly by how the publisher sets the reserve prices. In particular, because the publisher knows that it can sell the impression in Exchange 2 if it is left unsold in Exchange 1, the publisher sets a higher reserve price for Exchange 1 than if Exchange 2 did not exist. As such, the existence of an advertiser in Exchange 2 increases the expected payment of an advertiser in Exchange 1, and the existence of an advertiser in Exchange 1 lowers the probability of winning for an advertiser in Exchange 2. In the following lemma, we show how the optimal reserve prices are set by the publisher under both first-price and second-price auctions.

Lemma 2: Under waterfalling, regardless of the auction format (i.e., in both second-price and first-price auctions), the optimal reserve price of Exchange 2 is 1/2 and the optimal reserve price of Exchange 1 is

$$1 - \frac{1}{n_2 + 1} + \frac{1}{(n_2 + 1) \times 2^{n_2 + 1}}.$$

Lemma 2 has two interesting implications. First, it shows that the optimal reserve prices are not affected by the format of the auction. In fact, as we show in Proposition 1, the revenues of the publisher and both exchanges are the same in first-price auctions as in second-price auctions. The second implication of Lemma 2 is regarding the value of the optimal reserve prices. As we can see, the optimal reserve price of the second exchange is 1/2, regardless of the number of advertisers in each exchange. Under waterfalling, r_2 matters only when the impression is left unsold in Exchange 1. Therefore, if the publisher does not sell the impression in Exchange 2, it will generate zero revenue. This reduces the revenue maximization of the publisher for Exchange 2 to a standard optimal auction setting (with no "outside option" for the seller). In fact, the optimal reserve price 1/2 in Lemma 2 is Myerson's (1981) optimal reserve price for the case of uniform value distributions for the bidders.[10]

We can see from Lemma 2 that the optimal reserve price of Exchange 1 is always greater than or equal to 1/2, and only a function of n_2, the number of advertisers in Exchange 2. Intuitively, when the publisher is setting the reserve price of Exchange 1, it has to take its expected revenue from Exchange 2 into account. The expected revenue from Exchange 2 acts as an "outside option" for the publisher when selling its impression in Exchange 1 (i.e., the publisher's revenue if the impression is left unsold in Exchange 1). As n_2 increases, the publisher's expected revenue—and thus the value of "keeping" the impression—from Exchange 1 increases; therefore, the optimal reserve price of Exchange 1 increases as n_2 increases. In Proposition 1, we compare the revenue of the publisher and the exchanges under first-price auctions with their revenues under second-price auctions.

Proposition 1: Under waterfalling, the revenue of the publisher, the revenue of the exchanges, and the advertisers' utilities are not affected by the auction format.

Proposition 1 shows that, under waterfalling, if an exchange moves from second-price to first-price auction, or vice versa, the move does not affect the revenue of the publisher or either of the exchanges. This result follows from the revenue equivalence principle (e.g., Krishna 2009). Basically, if an exchange changes its auction format from second price to first price, advertisers that now have to pay what they bid (rather than the next-highest bid) shade their bids. While the exchange (and the publisher) make more revenue from a given set of bids, the amount by which the advertisers lower their bids cancels out the publisher's extra revenue from a given set of bids. This is in fact a general result about symmetric bidders and is not driven by our assumption about the number of advertisers or the advertisers' valuations being uniformly distributed.

Proposition 1 explains why exchanges did not move to first-price auctions under waterfalling. When RTB was introduced in 2009, exchanges used the second-price auction, an already-popular auction format in the context of online search advertising. Advertisers were familiar with second-price auctions, and the truthful nature of the auction made bidding strategies relatively simple. Because, under waterfalling, first-price auctions are equivalent to second-price in terms of expected equilibrium

[8] Note that $n_1 + n_2$ can be less than 4, as some advertisers may not join any of the exchanges.

[9] Bidding "truthfully" here means that the advertisers bid their true valuations (e.g., there is no bid shading). We borrow this term from the mechanism-design literature (e.g., see Krishna 2009). Note that even in mechanisms where advertisers do not bid truthfully (e.g., a first-price auction), the advertisers' true types may still be revealed in equilibrium.

[10] Myerson's (1981) optimal reserve price when the cumulative distribution function (CDF) and probability density function (PDF) of bidders' valuation are F and f, respectively, is $\phi^{-1}(0)$ where $\phi(x) = x - \frac{1-F(x)}{f(x)}$.

revenue, exchanges had no reason to abandon the simple and already-accepted second-price auctions. It was only in 2017, after header bidding became widely popular among publishers, that some exchanges started experimenting with first-price auctions and eventually moved to the first-price auction format. Indeed, in Proposition 4, we show that under header bidding, first-price auctions and second-price auctions are no longer equivalent.

Next, we analyze the advertisers' choices of exchanges, and the fees that the exchanges set in equilibrium. For any given fees f_1 and f_2, the advertisers' choices of exchanges can have multiple equilibria. In the following proposition, we show that the exchanges can charge the advertisers a positive fee in at least some equilibria of the game.[11]

Proposition 2: Under waterfalling, the exchanges can obtain positive buyer-side revenue in equilibrium (i.e., there are equilibria in which $n_1 f_1 + n_2 f_2 > 0$).

Proposition 2 shows that, under waterfalling, the exchanges can obtain positive revenue by charging advertisers a type of fee that is referred to as buyer-side fee in this industry. In fact, before the move to first-price auctions, exchanges had been charging, and in many cases increasing, their buyer-side fees. Interestingly, in Proposition 5, we show that under header bidding with first-price auctions, the exchanges' ability to obtain positive buyer-side revenue disappears—that is, both exchanges charge zero fees (i.e., $f_1 = f_2 = 0$ in *all* equilibria of the game). Intuitively, the reason that exchanges can set positive buyer-side fees in equilibrium under waterfalling is that the exchanges can differentiate in their offerings to the advertisers. In particular, as we demonstrate in the following example, the order of the exchanges in the waterfall sequence, and the number of advertisers that each exchange has, can make one exchange more attractive to the advertisers than the other.

Example 1: There is an equilibrium where the fees are $f_1^* = 0$, $f_2^* = \frac{n\times(4^n-1)-(2^n-1)^2-n^2\times 2^n}{n^2(n+1)\times 2^{2n+1}}$, and all advertisers join Exchange 2. Each advertiser's expected utility in this equilibrium is $\frac{1}{n(n+1)} - \frac{1}{n\times 2^n} + \frac{1}{(n+1)\times 2^{n+1}} - f_2^*$; the publisher's expected revenue is $(1-f)\times\left(1-\frac{2}{n+1}+\frac{1}{(n+1)\times 2^n}\right)$. The revenue of Exchange 1 is 0, and the expected revenue of Exchange 2 is $nf_2^* + f\times\left(1-\frac{2}{n+1}+\frac{1}{(n+1)\times 2^n}\right)$.[12]

In the equilibrium in Example 1, even though Exchange 1 has no fee, all advertisers join Exchange 2, which has a positive fee. Intuitively, if an advertiser switches to Exchange 1, it will face a very high reserve price because the publisher has a high expected revenue from Exchange 2 (as discussed in Lemma 2). As such, advertisers are better off paying fee f_2 and staying with Exchange 2 than deviating to Exchange 1. In other words, under waterfalling, exchanges with more advertisers are able to set a higher buyer-side fee.[13] As we show in the following subsection, this advantage disappears when exchanges move to header bidding with first-price auctions.

Before concluding this section, and as a segue into the discussion of header bidding, we present Proposition 3. This proposition examines an off-equilibrium path subgame solution that is meant to reflect the state of the industry just before header bidding was deployed by publishers and to model their consideration in switching from waterfalling to header bidding. It describes a specific situation in which exchanges use the second-price auction format that was widely popular and advertisers cannot move between exchanges for exogenous reasons such as inertia.

Proposition 3: Assuming that the advertisers have chosen their exchanges and the exchanges use second-price auctions, for any $n_1, n_2 > 0$, the revenue of the publisher when using header bidding is higher than when using waterfalling.

Proposition 3 shows that the publisher's revenue under header bidding is higher than under waterfalling, if the exchanges keep using second-price auctions. We emphasize, however, that this result is off the equilibrium path and applies to the subgame where advertisers have already joined exchanges (i.e., as we show in the following subsection, under header bidding, both exchanges use first-price auctions in equilibrium). Our goal here is to show that the publishers' move to header bidding was not necessarily motivated by the exchanges' subsequent move to first-price auctions. In fact, there is no evidence that publishers anticipated the adoption of first-price auctions by exchanges when they moved to header bidding, and, for almost a year after the publisher's move to header bidding, the exchanges continued to run second-price auctions. Proposition 3 explains the rapid growth of header bidding during a time when exchanges were still using second-price auctions.[14]

Header Bidding

In this section, we analyze the advertisers', the exchanges', and the publisher's strategies under header bidding. As in the preceding section, we use backward induction to solve the game. As before, we assume that n_1 and n_2 advertisers have joined Exchanges 1 and 2, respectively, and the reserve prices are r_1 and r_2. The following lemma summarizes the advertisers' bidding strategies in first-price and second-price auctions.

Lemma 3: Under header bidding, if an exchange uses a second-price auction, all advertisers within that exchange bid truthfully. If Exchange i uses a first-price auction, assuming

[11] In the Web Appendix, we prove that under some mild assumptions, the exchanges obtain positive total buyer-side revenue in all equilibria of the game.

[12] We can prove that this equilibrium has the highest expected publisher revenue and the highest total exchange revenue among all equilibria of the game.

[13] In the Web Appendix, we show a similar result when the decision for the ordering of the exchanges under waterfalling is endogenous in the model.

[14] We note that Sayedi (2018) has a similar result in a model with no exchanges and two horizontally differentiated advertisers.

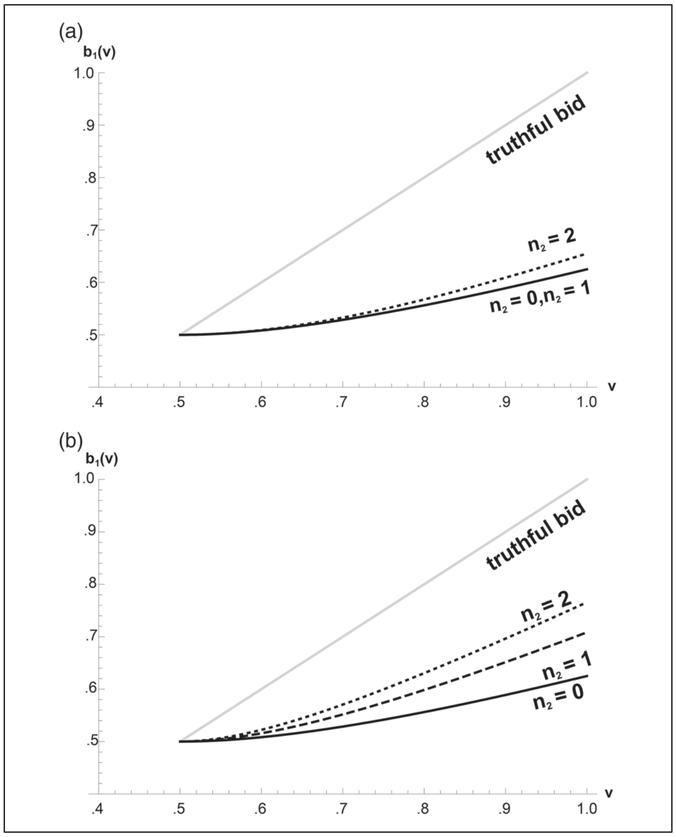

Figure 1. Bids of advertisers in Exchange 1 as a function of their valuation, for $n_1 = 2$, reserve prices $r_1 = r_2 = 1/2$, and different values of n_2.

Notes: The gray line (i.e., truthful bid) is for when Exchange 1 uses a second-price auction (for any value of n_2); the black curves are for when it uses a first-price auction.

that $r_1 = r_2 = 1/2$, the bid of an advertiser with valuation v who uses Exchange i is as follows.

- If both exchanges use first-price auctions: $v - \frac{v}{n_1+n_2} + \frac{1}{(n_1+n_2)\times 2^{n_1+n_2}\times v^{n_1+n_2-1}}$.
- If the other exchange uses a second-price auction, the advertisers' bid function, which varies depending on specific values of n_1 and n_2, is presented in Appendix A.

Note that we have stated the second part of Lemma 3, when at least one exchange uses a first-price auction, only for the case of $r_1 = r_2 = 1/2$. Although we can calculate the bidding strategies in more general cases (and they will give us more cumbersome expressions), it turns out that this is the only subgame that is needed for the analysis of the game (i.e., as we prove subsequently, the publisher always sets $r_1 = r_2 = 1/2$ in equilibrium, and other subgames are all dominated).

Lemma 3 shows that, unlike in Lemma 1, advertisers' strategies under header bidding in an exchange can directly depend on what happens in the other exchange. In particular, when an exchange uses a first-price auction, advertisers within that exchange take the existence of the other exchange, and the number of advertisers within the other exchange, into account when calculating their bids. Interestingly, if an exchange uses a second-price auction, advertisers within that exchange do not take the existence of the other exchange into account, as bidding truthfully continues to be a weakly dominant strategy. This is illustrated in Figure 1, where the gray line depicting the advertisers' bids in a second-price auction is not affected by the number of advertisers in the other exchange, n_2, whereas the advertisers' bids in a first-price auction increase as n_2 increases (i.e., the "outside competition" becomes stronger).

Figure 1 shows that advertisers submit lower bids in a first-price auction than in a second-price auction. However, the lower bids do not imply a lower revenue for the exchange because advertisers pay what they bid, instead of the next-highest bid, in a first-price auction. In fact, from the revenue equivalence principle (e.g., Krishna 2009) we know that the expected clearing price of a second-price auction is the same as the expected clearing price of a first-price auction when $n_2 = 0$. In other words, the truthful bidding function (i.e., the solid gray lines in Figure 1) in a second-price auction has the same expected clearing price as the solid black curve bidding function (i.e., cases of $n_2 = 0$) in a first-price auction. Therefore, it is easy to see that when $n_2 = 2$, the equilibrium bidding function, represented by the dotted black line, leads to a higher expected clearing price than a first-price auction with $n_2 = 0$, thus also a higher expected clearing price than a second-price auction (for any n_2).

Next, we discuss two important effects of first-price auctions under header bidding:

1. **Exposure to outside competition:** Under header bidding, when Exchange i uses a first-price auction, unlike in Lemma 1, advertisers take the existence of the other exchange into account when calculating their bids. This is because advertisers know that, under header bidding, just being the highest bidder in their own exchange is not sufficient for winning the impression. Note that this aggressive bidding behavior does not happen if Exchange i uses a second-price auction, where bidding truthfully continues to be a weakly dominant strategy. In other words, by using a first-price auction, an exchange can expose its advertisers to the outside competition, and, therefore, induce them to bid more aggressively. The exchange benefits from this exposure as it increases its expected clearing price as well as the probability of winning the impression. Note that first-price auctions allow exchanges to expose advertisers to outside competition only under header bidding. In particular, the exposure effect happens because of the *parallel* nature of header bidding (i.e., all advertisers, regardless of what exchange they are in, are competing for the impression simultaneously).
2. **Unified first-price auction:** If both exchanges use first-price auctions with the same reserve price,

advertisers' equilibrium bids are as if they are all in one unified first-price auction. In other words, the negative effect of the advertisers being in two separate markets (on the publisher's revenue) disappears.

Next, we analyze the exchanges' choice of first-price versus second-price auctions.

Proposition 4: Under header bidding, when $r_1 = r_2 = 1/2$ and for any values of $n_1, n_2 > 0$, there is a unique equilibrium where both exchanges use first-price auctions.

Intuitively, the result of Proposition 4 is driven by the outside-competition-exposure effect of first-price auctions under header bidding. In particular, we know from the revenue equivalence principle that if Exchange j did not exist, the revenue of Exchange i under first-price auction would have been the same as under second-price auction. In the presence of Exchange j, the bids in a second-price auction in Exchange i remain truthful, and thus, the expected clearing price of Exchange i remains unchanged. However, the outside-competition-exposure effect of first-price auctions, as discussed previously, induces the advertisers to bid more aggressively (compared with when Exchange j did not exist). As such, by switching to the first-price format, an exchange can increase its expected clearing price and, therefore, its expected revenue. This leads to the unique equilibrium of Proposition 4 where both exchanges use first-price auctions.

The result of Proposition 4 is conditioned on the publisher setting the reserve price $r_1 = r_2 = 1/2$ for both exchanges. Next, we show that $r_1 = r_2 = 1/2$ is indeed the optimal pair of reserve prices for the publisher. As discussed previously, we know when both exchanges use first-price auctions with the same reserve prices, advertisers' bids from both exchanges are as if they are all in one unified first-price auction. In other words, the equilibrium in Proposition 4 is equivalent to the equilibrium of a first-price auction with all $n_1 + n_2$ advertisers and reserve price 1/2. Most importantly, the reserve price 1/2 is already the optimal reserve price of a unified first-price auction[15] (i.e., the publisher can achieve the revenue of Myerson's [1981] optimal mechanism by setting $r_1 = r_2 = 1/2$). In more detail, suppose that M_1 is the mechanism that runs a simple first-price auction among the n advertisers with reserve price $r = \phi^{-1}(0) = 1/2$. Consider the family of mechanisms $M_2(r_1, r_2)$, where we split the advertisers into n_1 and n_2 among the two exchanges (for any choice of n_1 and n_2 summing to n), let the two exchanges decide if they want to use a first-price or a second-price auction for their own group of advertisers with reserve prices r_1 and r_2, respectively, take the two resulting clearing prices from each group, and pick the largest one to be the winner. Myerson's result implies that

$$\text{Revenue}(M_1) \geq \max_{r_1, r_2} \{\text{Revenue}[M_2(r_1, r_2)]\}.$$

Now let M_3 be the specific mechanism in the family $M_2(r_1, r_2)$, where both of the reserve prices are equal to $\phi^{-1}(0) = 1/2$. By our previous observation on unified first-price auctions, because both exchanges are using a first-price auction and both reserve prices are equal to $\phi^{-1}(0)$ in M_3, it is essentially identical to a unified first-price auction (i.e., the mechanism M_1). Moreover, under M_3, from Proposition 4 we know that in the subgame equilibrium between the two exchanges, both the exchanges will choose a first-price auction. Thus, we get

$$\text{Revenue}(M_3) = \text{Revenue}(M_1) \geq \max_{r_1, r_2} \{\text{Revenue}[M_2(r_1, r_2)]\}.$$

In particular, this shows that for all splits (n_1, n_2), the revenue of M_3 when setting $r_1 = r_2 = 1/2$ is at least as large as that for any pair of reserve prices (including cases when $r_1 \neq r_2$), and this revenue is the maximum possible over all ways of selling the impression. Thus, we get the following result.

Corollary 1: Under header bidding, the publisher can achieve Myerson's (1981) optimal revenue.

To highlight the significance of Corollary 1, note that Myerson's (1981) setting has far fewer constraints for the seller than our setting. In particular, under Myerson's setting, the seller has direct access to the bidders and full control over the selling mechanism. In our setting, there are intermediaries (i.e., exchanges) that choose mechanisms that optimize their own revenue. Yet we get a unique equilibrium where the intermediaries' strategies optimize the seller's revenue. Interestingly, the revenue of the exchanges is not necessarily optimized in this equilibrium (e.g., when $n_1 = 3$ and $n_2 = 1$), Exchange 1 is better off when both exchanges use second-price auctions than when they both use first-price.

Intuitively, each exchange deviates to a first-price auction to increase its expected clearing price and, therefore, increase its expected seller-side revenue. However, when both exchanges move to first-price auctions, each faces a higher expected clearing price from the other exchange. Therefore, an exchange may end up with a lower equilibrium revenue when both exchanges move to first-price auctions than when they both use second-price auctions. This resembles (even though it is not mathematically equivalent to) a prisoner's dilemma situation for exchanges with respect to their choices of their auction format. As we show in Proposition 5, the negative effect of moving to first-price auctions for the exchanges becomes even stronger when we take buyer-side fees into account.

Proposition 4 provides a new explanation for why ad exchanges moved from second-price to first-price auctions. Our explanation is consistent with the timing of the transition to first-price auctions as if it were triggered by the publisher's adoption of header bidding and is consistent with how this

[15] Myerson's (1981) optimal reserve price, under both second-price and first-price auctions, when the CDF and PDF of bidders' valuation are F and f, respectively, is $\phi^{-1}(0)$, where $\phi(x) = x - \frac{1-F(x)}{f(x)}$ is Myerson's "virtual valuation." When F is uniform [0, 1], the optimal reserve price becomes 1/2.

market evolved over time. The market share of header bidding among the top publishers in the United States grew from 0% to over 70% from 2014 to 2016 (Insider Intelligence 2016). Exchanges started experimenting with first-price auctions in 2017, and all major exchanges fully moved to first-price auction by early 2019 (Sluis 2017a).

A disparity in how major platforms such as Google and Facebook sell display advertising impressions is also consistent with our explanation. Google only uses first-price auctions in its exchange platform, where it is an intermediary. For selling its own inventory, such as impressions on YouTube, however, because header bidding is not used (i.e., advertisers have to purchase the impressions directly from Google), Google still uses second-price auctions (Sluis 2019). Similarly, Facebook uses a generalized version of second-price auctions to sell its display advertising impressions directly to advertisers. The fact that these large firms have adopted first-price auctions only in situations where they act as intermediaries provides additional support for our explanation.

Next, we continue our analysis of the game. Note that because the publisher can achieve the optimal revenue by setting $r_1 = r_2 = 1/2$, we do not have to solve for other subgames with other values of r_i. Thus, we move to advertisers' equilibrium strategies regarding which exchange to join and exchanges' equilibrium fees under header bidding. The following proposition summarizes the exchanges' equilibrium fees, and advertisers' choice of exchanges under header bidding.

Proposition 5: Under header bidding, both exchanges set their buyer-side fees to zero (i.e., $f_1 = f_2 = 0$), in equilibrium. The advertisers are indifferent about which exchange to join.

From Proposition 4, we know that both exchanges use first-price auctions in equilibrium. This reduces the market with two exchanges to a unified first-price auction. Because advertisers are forward-looking, they know that the choice of exchange does not affect their probability of winning or their expected utility. In other words, regardless of the choice of the exchange, an advertiser wins if and only if he has the highest bid among all advertisers across all exchanges. Therefore, for any buyer-side fees f_1 and f_2, an advertiser's optimal strategy is to choose the exchange with the lower fee. Given the advertisers' strategies, the exchanges set their buyer-side fees to zero in equilibrium. The finding of Proposition 5 is in line with industry reports that show, after adoption of first-price auctions, many exchanges have reduced, or completely dropped, buyer-side fees (Sluis 2017c).

Note that, using Proposition 2, we know that exchanges can extract positive equilibrium revenue through buyer-side fees when the publisher uses waterfalling. However, Proposition 5 shows that, when the publisher uses header bidding, the exchange's ability to charge buyer-side fees disappears. Intuitively, this is because the exchanges can differentiate in terms of their position in the waterfall sequence and the number of advertisers they have under waterfalling. For example, if an exchange has more advertisers on its platform, or if it has a more favorable position in the waterfall sequence, it can charge higher buyer-side fees in equilibrium. However, when the publisher uses header bidding and the exchanges use first-price auctions, the exchanges' ability to differentiate themselves using their position in the sales channel disappears. Advertisers choose the exchange with the lowest fee, and the exchanges' profit from buyer-side fees declines to zero.

Finally, we should mention that while the exchanges' buyer-side fees decline to zero in equilibrium, exchanges still obtain positive revenue through seller-side fees. In other words, the exchanges still benefit from attracting more advertisers because, by having more advertisers, they can increase their seller-side revenue. We should also note that, in our model, the exchanges' buyer-side fees decline all the way to zero because, to facilitate exposition, we have assumed that the exchanges are ex ante identical (no differentiation). In practice, exchanges can horizontally differentiate though the tools and services that they offer to the advertisers. Indeed, in Appendix B, we consider an extension of our main model with horizontally differentiated exchanges and show that the exchanges obtain positive buyer-side revenue in equilibrium when they are horizontally differentiated.

Conclusion

In this article, we propose a simple model of RTB in display advertising to analyze the evolution of selling mechanisms in this market and the consequences for advertisers, publishers, and exchange platforms. We show that, when exchanges were using second-price auctions and advertisers' exchange affiliations were unchanged, the publisher's revenue when using header bidding is always higher than waterfalling; this result explains the rapid adoption of header bidding by publishers in recent years.

Our results also provide a new explanation for why exchange platforms moved from second-price auctions to first-price. Second-price auctions have been the industry standard in online advertising for over a decade. In fact, they are still the dominant selling mechanism in search advertising and display advertising markets in which publishers directly sell to advertisers (without going through a third-party exchange). For example, Google uses second-price auctions for selling YouTube and AdSense impressions, and Facebook uses a generalized form of second-price auctions to sell its display advertising inventory. We argue that this move may be due to the wide adoption of header bidding. Our results provide managerial implications for advertisers, publishers, and exchanges in the online advertising industry.

Implications for Advertisers

In the past few years, the selling mechanism in RTB market has dramatically changed: first, publishers moved from waterfalling to header bidding, and then exchanges moved from second-price to first-price auctions. This leaves the advertisers uncertain about how to adjust their bidding strategies under the new mechanism. Our results show that advertisers should

shade their bids using the same methods as in a standard first-price auction. The degree of shading depends on the number of other advertisers in the market as well as their distribution of values for the impression. Under header bidding, in contrast with waterfalling, advertisers should consider every other advertiser in the market as competition—not only those who use the same exchange.

Another important implication for advertisers is regarding the choice of exchange. Previously, under waterfalling, advertisers had to pay attention to the position of an exchange in the sequence of the waterfall, as well as the number of other advertisers in each exchange. Using the exchange with the lowest fee was not necessarily the optimal strategy. This continues to hold under header bidding with second-price auctions. However, under header bidding with first-price auctions, submitting the bid through different exchanges does not affect the final price and allocation of an advertiser; as such, the optimal strategy of an advertiser is the simple one of using the exchange with the lowest fee.

Implications for Exchanges

Under waterfalling, an exchange could use its position in the waterfall sequence to differentiate itself from other exchanges. When exchanges use second-price auctions, they can use their set of the advertisers to differentiate themselves from other exchanges; intuitively, an advertiser benefits from being in an exchange with other advertisers with similar valuations. However, the combination of header bidding and first-price auctions puts exchanges in direct competition. While the move to first-price auction was necessary for an exchange to survive in the short run after the publishers adopted header bidding, after taking its effect on advertisers' choices of exchanges into account, our results show that the move will lower the exchanges' equilibrium buyer-side fees in the long run. This is consistent with several industry reports indicating a steep decline in exchange fees since the adoption of first-price auctions (Sluis 2017a).

To avoid head-on competition, exchanges can no longer rely on their position within the selling mechanism as a point of differentiation; they have to create new strategies to differentiate in this market. For example, they can leverage their information about transactions in this market and offer analytical tools to advertisers that use their platform. Alternatively, they can vertically differentiate by filtering out low-quality (or suspicious/fraudulent) impressions to guarantee certain viewability rates.

Implications for Publishers

Our results indicate that publishers mainly benefit from the adoption of first-price auctions by exchanges. The direct benefit of first-price auctions for the publishers is that all advertisers directly compete with each other when all exchanges use first-price auctions. In other words, even though the advertisers only compete within an exchange, they take the existence of other exchanges (and other advertisers in those exchanges) into account when optimizing their bids. This effectively becomes equivalent to as if all advertisers were in the same unified first-price auction. In other words, first-price auctions under header bidding eliminate the negative effect of advertisers being separated into multiple exchanges on the publisher's revenue. There is already some early evidence of this improved revenue for publishers and a more competitive market for advertisers as a result of the move to first-price auctions in the industry (Bigler 2019).

We also show that first-price auctions under header bidding allow publishers to achieve the revenue of Myerson's (1981) optimal mechanism. In other words, just by setting the reserve prices optimally, the publisher can achieve the revenue of the optimal mechanism that has direct access to advertisers (i.e., without the interference of intermediaries) and can use any (individually rational) pricing and allocation.

Finally, from a computational point of view, the move to first-price auction simplifies the publisher's choice of optimal reserve prices. Under waterfalling, and also when exchanges use second-price auctions, the publisher has to solve a joint optimization problem and set asymmetric reserve prices even when the exchanges and the advertisers are symmetric. The reserve price of an exchange had to take the expected revenue from other exchanges into account. These interactions made the calculation of the optimal reserve prices extremely complicated. However, under header bidding with first-price auctions, the optimal reserve price of each exchange is independent of all other exchanges and happens to be Myerson's (1981) standard optimal reserve price.

Limitations and Future Research

Our results shed light on the evolution of the selling mechanism in the RTB market over time and provide insights for managers in this industry on how to buy and sell display advertising inventory in the current market. In our model, we make several simplifying assumptions. First and foremost, we assume that the advertisers' valuations are independent draws from the same distribution. In reality, advertisers may be asymmetric in their distributions of valuations; furthermore, their valuations may be correlated. Intuitively, Zeithammer (2019) shows that asymmetric distributions can favor adoption of first-price auctions, whereas Milgrom and Weber (1982) show that correlation in advertisers' valuations favors adoption of second-price auctions. The question of the optimal mechanism design in display advertising markets with general distributions is beyond the scope of this article and, thus, left to future research.

Similarly, our model assumes that the number of advertisers in the market is common knowledge. While this might hold in mature markets where advertisers have been competing with each other for a long time, it does not hold in other situations. In particular, an advertiser may not know how many other advertisers may be interested in a given impression. This assumption significantly simplifies our analysis and allows us

to use the framework in Myerson (1981). Studying advertising auctions where the number of competing advertisers is not known is an interesting research direction that could be explored in future research.

Our work is among a small set of articles that study the fast-growing RTB market. While we focused on the role of exchanges, future research could explore other entities in this market such as demand-side platforms, supply-side platforms, and data-management platforms. In particular, because competing advertisers sometimes use the same demand-side platforms, and competing publishers can use the same supply-side platforms, extending our analysis of intermediaries to such buyer and seller agents could lead to further interesting economic insights. Another direction to explore is how privacy regulation affects the role of data-management platforms and the optimal level of information sharing in this market.

Appendix A: Proofs and Alternative Timelines

Analyses and Proofs

Proof of Lemma 1. If Exchange i is using a second-price auction, it is easy to verify that bidding truthfully is a weakly dominant strategy for the advertisers.

Suppose that Exchange i is using a first-price auction, and let $b_i(v)$ be a symmetric increasing bidding function of advertisers in Exchange i. The function $b_i(v)$ should satisfy the boundary condition $b_i(r_i) = r_i$. The expected utility of an advertiser with valuation $v \geq r_i$ if it bids $b_i(x)$ instead of $b_i(v)$, for some $x \geq r_i$, is

$$u(x) = F(x)^{n_i-1}[v - b_i(x)] = x^{n_i-1}[v - b_i(x)].$$

To have an equilibrium, u(x) must be maximized for $x = v$. We have $\partial u/\partial x = (n_i - 1)x^{n_i-2}[v - b_i(x)] - x^{n_i-1}b_{i'}(x)$. Because u(x) is maximized for $x = v$, we get $\partial u/\partial x|_{x=v} = 0$, which gives the differential equation

$$b_{i'}(v) = \frac{(n_i - 1)[v - b_i(v)]}{v}.$$

The solution to this differential equation is $b_i(v) = v - \frac{v}{n_i} + \frac{C}{v^{n_i-1}}$, where C is a constant. Using the boundary condition $b_i(r_i) = r_i$, we can find that $C = \frac{r_i^{n_i}}{n_i}$. Therefore, in equilibrium, an advertiser in Exchange i with valuation $v \geq r_i$ bids $b_i(v) = v - \frac{v}{n_i} + \frac{r_i^{n_i}}{n_i \cdot v^{n_i-1}}$.

Proof of Lemma 2. From Myerson (1981), we know that the virtual valuations in Exchange 2 are defined as $\phi(x) = x - \frac{1-F(x)}{F'(x)} = 2x - 1$. Therefore, the optimal reserve price for Exchange 2, regardless of the auction format (according to Myerson), is $r_2 = \phi^{-1}(0) = 1/2$. Consequently, the expected revenue of the publisher from Exchange 2 is

$$\begin{aligned} w_2 &= (1-f) \times \left(n_2 \left\{ r_2[1 - F(r_2)]F(r_2)^{n_2-1} + \int_{r_2}^{1} y[1 - F(y)](n_2 - 1)F(y)^{n_2-2}F'(y)\,dy \right\} \right) \\ &= (1-f) \times \left[1 - \frac{2}{n_2+1} + \frac{1}{(n_2+1) \times 2^{n_2}} \right]. \end{aligned}$$

The expected revenue of the publisher from Exchange 2 is the publisher's expected revenue from not selling the impression in Exchange 1 (i.e., when analyzing Exchange 1, w_2 is the seller's value for keeping the item). Therefore, we know from Myerson (1981) that the optimal reserve price for Exchange 1 is $r_1 = \phi^{-1}\left(\frac{w_2}{1-f}\right) = 1 - \frac{1}{n_2+1} + \frac{1}{(n_2+1) \times 2^{n_2+1}}$. Note that, using Myerson, this is the optimal reserve price for both first-price and second-price auction formats.

Proof of Proposition 1. This result comes from the revenue equivalence principle. Suppose that Exchange i has n_i advertisers and it is using a second-price auction with reserve price r_i. Then its expected clearing price is

$$p_i^{SP} = n_i \left\{ r_i[1 - F(r_i)]F(r_i)^{n_i-1} + \int_{r_i}^{1} y[1 - F(y)](n_i - 1)F(y)^{n_i-2}F'(y)\,dy \right\}.$$

Similarly, suppose that Exchange i has n_i advertisers and is using a first-price auction with reserve price r_i. Let $b_i(v)$ be the bidding function of its advertisers. Then, its expected clearing price is

$$p_i^{FP} = n_i \int_{r_i}^{1} b_i(y)F(y)^{n_i-1}F'(y)\,dy.$$

It is easy to verify that for $b_i(v) = v - \frac{v}{n_i} + \frac{r_i^{n_i}}{n_i \times v^{n_i-1}}$, which is the bidding function of advertisers in a first-price auction (Lemma 1), it holds that $p_i^{SP} = p_i^{FP}$. Therefore, for a fixed reserve price, exchanges are indifferent between running a first-price auction or a second-price auction, as both will give the same expected revenue. When the publisher sets the reserve prices, the optimal reserve prices r_1 and r_2 are the same in a first-price auction as in a second-price auction (by Lemma 2). Therefore, the publisher's and the exchanges' expected revenues are the same in a first-price auction as in a second-price auction.

Finally, an advertiser's winning probability and expected payment are the same in a first-price auction and in a second-price auction. The winning probability is the probability that the advertiser has the highest valuation among the advertisers in his exchange, and the expected payments are simply p_i^{SP}/n_i and p_i^{FP}/n_i, for a second-price and a first-price auction respectively, which are equal. Therefore advertisers' utilities are not affected by the auction format.

Proof of Proposition 2. If Exchange 1 has n_1 advertisers and Exchange 2 has n_2 advertisers, then the expected utility of an

advertiser in Exchange 1 (ignoring the fee for now) is

$$u_1(n_1, n_2) = \int_{r_1}^{1} F(y)^{n_1-1} yF'(y)\,dy - \Big\{ r_1[1 - F(r_1)]F(r_1)^{n_1-1}. + \int_{r_1}^{1} y[1 - F(y)](n_1 - 1)F(y)^{n_1-2}F'(y)\,dy \Big\} = \frac{1}{n_1(n_1+1)} - \frac{r_1^{n_1}}{n_1} + \frac{r_1^{n_1+1}}{n_1+1}. \quad \text{(A1)}$$

The expected utility of an advertiser in Exchange 2 (again ignoring the fee) is

$$u_2(n_1, n_2) = F(r_1)^{n_1} \times \Big[\int_{r_2}^{1} F(y)^{n_2-1} yF'(y)\,dy - \times \Big\{ r_2[1 - F(r_2)]F(r_2)^{n_2-1}. + \int_{r_2}^{1} y[1 - F(y)](n_2 - 1)F(y)^{n_2-2}F'(y)\,dy \Big\} \Big] = r_1^{n_1} \times \Big[\frac{1}{n_2(n_2+1)} - \frac{r_2^{n_2}}{n_2} + \frac{r_2^{n_2+1}}{n_2+1} \Big]. \quad \text{(A2)}$$

Note that in Equations A1 and A2, r_1 and r_2 are the optimal reserve prices from Lemma 2 (i.e., r_1 is a function of n_2 and $r_2 = 1/2$).

Consider the case where $n_1 = 0$, $n_2 = n$, and $f_1 = 0$ (i.e., all advertisers are in Exchange 2 while Exchange 1 has a zero fee). Then, Exchange 2 can charge a positive fee $f_2 > 0$ such that no advertiser benefits by moving to Exchange 1 (even though Exchange 1's fee is 0). More specifically, the maximum fee Exchange 2 can charge so that none of its advertisers wants to move is

$$f_2^* = u_2(0, n) - u_1(1, n-1) = \frac{n \times (4^n - 1) - (2^n - 1)^2 - n^2 \times 2^n}{n^2(n+1) \times 2^{2n+1}} > 0.$$

In other words, there is an equilibrium where the total exchange buyer-side revenue is positive (i.e., $n_1f_1 + n_2f_2 > 0$).

Proof of Proposition 3. Recall that we are analyzing the situation where the number of advertisers in both exchanges is fixed and nonzero and both exchanges are using second-price auctions. First, note that it is weakly dominant for advertisers to bid truthfully under both waterfalling and header bidding. Therefore, moving to header bidding does not affect advertisers' bidding strategies. Moreover, for any pair of reserve prices r_1 and r_2, the publisher's revenue under header bidding is greater than or equal to that under waterfalling. This is because, under header bidding, the publisher can see the clearing prices of both exchanges before deciding which exchange wins the impression, whereas under waterfalling, the publisher has to accept or decline the clearing price of Exchange 1 before seeing the clearing price of Exchange 2. Therefore, assuming that r_1^* and r_2^* are the optimal reserve prices under waterfalling, we know that the publisher's revenue when using r_1^* and r_2^* under header bidding is greater than when using r_1^* and r_2^* under waterfalling. As such, the publisher's optimal revenue under header bidding is greater than the publisher's optimal revenue under waterfalling.

Proof of Lemma 3. If an exchange is using a second-price auction, then it is a weakly dominant strategy for its advertisers to bid their true valuation. If both exchanges use first-price auctions with the same reserve price, then for the advertisers this is equivalent to a global first-price auction. Therefore, their bidding function is the standard bidding function for a first-price auction with $n_1 + n_2$ advertisers.

Consider the case where one of the exchanges uses a first-price auction and the other uses a second-price auction. Without loss of generality, suppose that Exchange 1 is using a first-price auction and Exchange 2 is using a second-price auction.

Consider a symmetric equilibrium bidding function $b(v)$ for the advertisers in Exchange 1. Let $r = 1/2$. The expected utility of an advertiser with valuation $v \geq r$ if his bid is $b(x)$ instead of $b(v)$ is

$$u(x) = \{n_2F[b(x)]^{n_2-1} - (n_2 - 1)F[b(x)]^{n_2}\}F(x)^{n_1-1}[v - b(x)]. \quad \text{(A3)}$$

To have an equilibrium, this function must be maximized for $x = v$.

We start with the simple case where $n_2 \leq 1$ and $n_1 \geq 2$. The utility function becomes $u(x) = F(x)^{n_1-1}[v - b(x)]$. Therefore, advertisers in Exchange 1 can ignore Exchange 2 and bid as if they are in a simple first-price auction with n_1 advertisers. In this case, the bidding function is

$$b(v) = \begin{cases} 0 & , \text{ if } v < r, \\ v - \dfrac{v}{n_1} + \dfrac{r^{n_1}}{n_1 \times v^{n_1-1}} & , \text{ if } v \geq r. \end{cases}$$

Next, consider the case where $n_1 = 1$ and $n_2 \geq 2$. The utility function becomes

$$u(x) = \{n_2F[b(x)]^{n_2-1} - (n_2 - 1)F[b(x)]^{n_2}\}[v - b(x)].$$

Let $z(v)$ be the point y that maximizes the function $u[b^{-1}(y)]$. For vs that satisfy $z(v) \geq r$, we get $b(v) = z(v)$. For vs that satisfy $z(v) < r$, we have $b(v) = r + \varepsilon$, where $0 < \varepsilon < v - r$.[16] Let t be such that $z(t) = r$. Then the bidding function is

$$b(v) = \begin{cases} 0 & , \text{ if } v \leq r, \\ r + \varepsilon & , \text{ if } r < v < t, \\ z(v) & , \text{ if } v \geq t. \end{cases}$$

From this, we can obtain the following special cases.

- For $n_1 = 1$ and $n_2 = 2$, it is $b(v) = \begin{cases} 0 & , \text{ if } v \leq r, \\ r + \varepsilon & , \text{ if } v > r. \end{cases}$

[16] The role of ε here is to break the tie between the two exchanges in the case that there is only one advertiser in Exchange 2 that bids above the reserve price. If we change the tie-breaking rule to say that the exchange with the first-price auction wins in a tie, then we can remove ε from the equations.

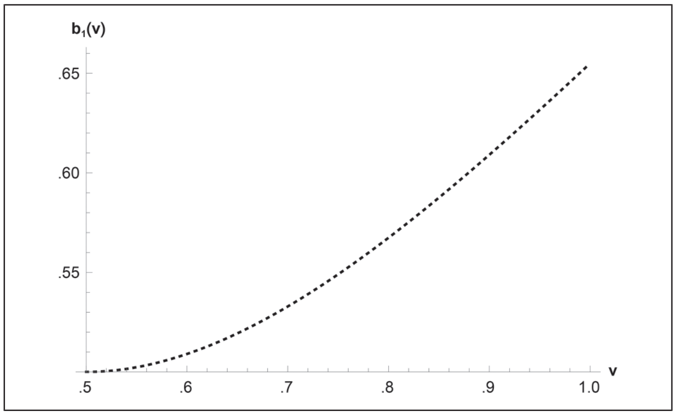

Figure A1. Bidding function of advertisers in exchange 1 for $n_1 = 2$ and $n_2 = 2$, when Exchange 1 is using a first-price auction and Exchange 2 is using a second-price auction.

- For $n_1 = 1$ and $n_2 = 3$, it is $b(v) = \begin{cases} 0 & , \text{ if } v \le r, \\ r+\varepsilon & , \text{ if } r < v < \frac{5}{6}, \\ \frac{1}{16}\left(6v + 9 - \sqrt{36v^2 - 84v + 81}\right) & , \text{ if } v \ge \frac{5}{6}. \end{cases}$

Some other trivial cases are the following.

- For $n_1 = 1$ and $n_2 = 1$, it is $b(v) = \begin{cases} 0 & , \text{ if } v \le r, \\ r+\varepsilon & , \text{ if } v > r. \end{cases}$
- For $n_1 = 1$ and $n_2 = 0$, it is $b(v) = \begin{cases} 0 & , \text{ if } v < r, \\ r & , \text{ if } v \ge r. \end{cases}$

Finally, consider the case with $n_1 = 2$ and $n_2 = 2$. To find b, we need to solve the differential equation $\partial u / \partial x|_{x=v} = 0$ with boundary condition $b(r) = r$. The differential equation is

$$b'(v) = \frac{[2-b(v)]b(v)[v-b(v)]}{v[-3b(v)^2 + 2(v+2)b(v) - 2v]}, \tag{A4}$$

and its solution is plotted in Figure A1. Note that even without a closed-form solution for b(v) for the case of $n_1 = n_2 = 2$, we can prove Proposition 4 analytically by using an analytical lower bound for b(v).

Proof of Proposition 4. If both exchanges use a second-price auction, the expected seller-side revenue of Exchange 1 is

$$\begin{aligned} L_1^{(SP,SP)} = {} & f \times \Big\{r_1 \times n_1[1-F(r_1)]F(r_1)^{n_1-1}F(r_2)^{n_2} + \frac{1}{2}\times r_2 \times n_1 n_2[1-F(r_1)]F(r_1)^{n_1-1}[1-F(r_2)]F(r_2)^{n_2-1} \\ & + F(r_2)^{n_2}\left\{\int_{r_1}^{a} y \times n_1(n_1-1)[1-F(y)]F(y)^{n_1-2}F'(y)\,dy\right\} + n_2[1-F(r_2)]F(r_2)^{n_2-1}\left\{\int_{r_2}^{a} y \times n_1(n_1-1)[1-F(y)]F(y)^{n_1-2}F'(y)\,dy\right\} \\ & + \int_{r_2}^{a}\left(n_2(n_2-1)[1-F(z)]F(z)^{n_2-2}F'(z)\left\{\int_{z}^{a} y \times n_1(n_1-1)[1-F(y)]F(y)^{n_1-2}F'(y)\,dy\right\}\right)dz\Big\}, \end{aligned}$$

where $a = 1$ is the maximum possible valuation of an advertiser. The expected seller-side revenue of Exchange 2 is

$$\begin{aligned} L_2^{(SP,SP)} = {} & f \times \Big\{r_2 \times F(r_1)^{n_1} n_2[1-F(r_2)]F(r_2)^{n_2-1} + \frac{1}{2}\times r_2 \times n_1 n_2[1-F(r_1)]F(r_1)^{n_1-1}[1-F(r_2)]F(r_2)^{n_2-1} \\ & + F(r_1)^{n_1}\left\{\int_{r_2}^{a} y \cdot n_2(n_2-1)[1-F(y)]F(y)^{n_2-2}F'(y)\,dy\right\} + n_1[1-F(r_1)]F(r_1)^{n_1-1}\left\{\int_{r_2}^{a} y \times n_2(n_2-1)[1-F(y)]F(y)^{n_2-2}F'(y)\,dy\right\} \\ & + \int_{r_2}^{a}\left[n_2(n_2-1)[1-F(z)]F(z)^{n_2-2}F'(z)\left\{\int_{r_1}^{z} z \times n_1(n_1-1)[1-F(y)]F(y)^{n_1-2}F'(y)\,dy\right\}\right]dz\Big\}. \end{aligned}$$

If both exchanges use a first-price auction, the expected seller-side revenue of Exchange 1 is

$$L_1^{(FP,FP)} = f \times \left\{r_1 n_1 F(r_1)^{n_1-1}[1-F(r_1)]F(r_2)^{n_2} + \int_{r_1}^{a} y(n_1+n_2-1)n_1[1-F(y)]F(y)^{n_1+n_2-2}F'(y)\,dy\right\}.$$

The expected seller-side revenue of Exchange 2 is

$$L_2^{(FP,FP)} = f \times \left\{r_2 n_2 F(r_2)^{n_2-1}[1-F(r_2)]F(r_1)^{n_1} + \int_{r_2}^{a} y(n_1+n_2-1)n_2[1-F(y)]F(y)^{n_1+n_2-2}F'(y)\,dy\right\}.$$

Next, suppose that Exchange 1 is using a first-price auction, while Exchange 2 is using a second-price auction. Let $b_1(v)$ be the bidding function of an advertiser in Exchange 1.[17] Then, the expected seller-side revenue of Exchange 1 is

$$L_1^{(FP,SP)}(b_1) = f \times \Bigg\{F(r)^{n_2}\int_r^a b_1(x)n_1F(x)^{n_1-1}F'(x)\,dx + n_2F(r)^{n_2-1}[1-F(r)]\int_r^a b_1(x)n_1F(x)^{n_1-1}F'(x)\,dx$$
$$+\int_r^a\left[b_1(z)n_1F(z)^{n_1-1}F'(z)\left\{\int_r^{b_1(z)}(n_2-1)n_2[1-F(x)]F(x)^{n_2-2}F'(x)\,dx\right\}\right]dz\Bigg\},$$

where $r = 1/2$ is the common reserve price. The expected seller-side revenue of Exchange 2 is

$$L_2^{(FP,SP)}(b_1) = f \times \Bigg\{r \times n_2F(r)^{n_2-1}[1-F(r)]F(r)^{n_1} + F(r)^{n_1}\int_r^a x(n_2-1)n_2[1-F(x)]F(x)^{n_2-2}F'(x)\,dx$$
$$+\int_r^a n_1F(z)^{n_1-1}F'(z)\left\{\int_{b_1(z)}^a x(n_2-1)n_2[1-F(x)]F(x)^{n_2-2}F'(x)\,dx\right\}dz\Bigg\}.$$

Finally, suppose that Exchange 2 is using a first-price auction, while Exchange 1 is using a second-price auction. Let $b_2(v)$ be the bidding function of an advertiser in Exchange 2. Then, the expected seller-side revenue of Exchange 1 is

$$L_1^{(SP,FP)}(b_2) = f \times \Bigg\{r \times n_1F(r)^{n_1-1}[1-F(r)]F(r)^{n_2} + F(r)^{n_2}\int_r^a x(n_1-1)n_1[1-F(x)]F(x)^{n_1-2}F'(x)\,dx$$
$$+\int_r^a n_2F(z)^{n_2-1}F'(z)\left\{\int_{b_2(z)}^a x(n_1-1)n_1[1-F(x)]F(x)^{n_1-2}F'(x)\,dx\right\}dz\Bigg\}.$$

The expected seller-side revenue of Exchange 2 is

$$L_2^{(SP,FP)}(b_2) = f \times \Bigg\{F(r)^{n_1}\int_r^a b_2(x)n_2F(x)^{n_2-1}F'(x)\,dx + n_1F(r)^{n_1-1}[1-F(r)]\int_r^a b_2(x)n_2F(x)^{n_2-1}F'(x)\,dx$$
$$+\int_r^a\left[b_2(z)n_2F(z)^{n_2-1}F'(z)\left\{\int_r^{b_2(z)}(n_1-1)n_1[1-F(x)]F(x)^{n_1-2}F'(x)\,dx\right\}\right]dz\Bigg\}.$$

Consider the following payoff matrix, $M_{n_1,n_2}(b_1, b_2)$, of the game between the exchanges.

		Exchange 2	
		SP	FP
Exchange 1	SP	$\{L_1^{(SP,SP)}, L_2^{(SP,SP)}\}$	$\{L_1^{(SP,FP)}(b_2), L_2^{(SP,FP)}(b_2)\}$
	FP	$\{L_1^{(FP,SP)}(b_1), L_2^{(FP,SP)}(b_1)\}$	$\{L_1^{(FP,FP)}, L_2^{(FP,FP)}\}$

To show that (FP, FP) is the unique equilibrium, it is sufficient to show that

$$L_1^{(FP,FP)} > L_1^{(SP,FP)}(b_2) \text{ and } L_2^{(FP,FP)} > L_2^{(FP,SP)}(b_1) \text{ and } L_1^{(FP,SP)}(b_1) > L_1^{(SP,SP)}.$$

Consider a pointwise lower bound $b_{1'}(v)$ of the bidding function $b_1(v)$. Note that $L_1^{(FP,SP)}(b_1) \geq L_1^{(FP,SP)}(b_{1'})$ and $L_2^{(FP,SP)}(b_1) \leq L_2^{(FP,SP)}(b_{1'})$. This is because if the advertisers in Exchange 1 decrease their bids, Exchange 1's revenue will go down (lower chance of winning and lower clearing price) while Exchange 2's revenue will go up (higher chance of winning).

Similarly, if $b_{2'}(v)$ is a pointwise lower bound of $b_2(v)$, it holds that $L_1^{(SP,FP)}(b_2) \leq L_1^{(SP,FP)}(b_{2'})$.

Therefore, it is sufficient to show that

$$L_1^{(FP,FP)} > L_1^{(SP,FP)}(b_{2'}) \text{ and } L_2^{(FP,FP)} > L_2^{(FP,SP)}(b_{1'}) \text{and} L_1^{(FP,SP)}(b_{1'}) > L_1^{(SP,SP)}. \quad \text{(A6)}$$

[17] Here, we assume a symmetric equilibrium bidding strategy for advertisers in Exchange 1.

for some lower-bound functions $b_{1'}$ and $b_{2'}$.

Consider the function

$$b_{1'}(v) = v - \frac{v}{n_1} + \frac{1}{n_1 \times 2^{n_1} \times v^{n_1 - 1}}.$$

This is how advertisers in Exchange 1 would bid if Exchange 1 were running a first-price auction and advertisers were completely ignoring the existence of Exchange 2. When the advertisers consider Exchange 2, their bids in equilibrium can only increase, because now they have to compete with a larger outside option. Therefore, $b_{1'}$ is a pointwise lower bound of b_1.

Similarly, the function $b_{2'}(v) = v - \frac{v}{n_2} + \frac{1}{n_2 \times 2^{n_2} \times v^{n_2 - 1}}$ is a pointwise lower bound of b_2.

For these lower bounds, the inequalities in Equation A6 become simple inequalities that involve only n_1 and n_2. Therefore, for given n_1 and n_2, it is easy to verify them.

Next, we consider all the cases for $n_2 \geq n_1 > 0$ and $n_1 + n_2 \leq n$.[18]

- For $n_1 = 2$ and $n_2 = 2$, the matrix $M_{2,2}$ for the lower-bound bidding functions $b_{1'}$ and $b_{2'}$ is

$$M_{2,2}(b_{1'}, b_{2'}) = f \times$$

{.272917, .272917}	{.208547, .334728}
{.334728, .208547}	{.30625, .30625}

- For $n_1 = 1$ and $n_2 = 3$, it is

$$M_{1,3}(b_{1'}, b_{2'}) = f \times$$

{.078125, .484375}	{.03125, .53125}
{.125, .4375}	{.153125, .459375}

- For $n_1 = 1$ and $n_2 = 2$, it is

$$M_{1,2}(b_{1'}, b_{2'}) = f \times$$

{.125, .354167}	{.0625, .416667}
{.1875, .291667}	{.177083, .354167}

- For $n_1 = 1$ and $n_2 = 1$, it is

$$M_{1,1}(b_{1'}, b_{2'}) = f \times$$

{.1875, .1875}	{.125, .25}
{.25, .125}	{.208333, .208333}

[18] The cases with $n_1 \geq n_2$ are symmetric.

We can see that for all cases, the inequalities in Equation A6 are satisfied. Therefore, for all cases, (FP, FP) is the unique equilibrium.

Proof of Proposition 5. Because advertisers are forward looking, at the time of choosing between the exchanges, they know that the reserve prices will be set at $r_1 = r_2 = 1/2$ and both exchanges will use first-price auctions. Intuitively, this implies that the exchanges are in a Bertrand competition when setting their buyer-side fees to attract advertisers. In the following, we formalize this intuition.

If $f_1 \neq f_2$, it is optimal for an advertiser to choose the exchange with the lower fee. Therefore, in any pure-strategy Nash equilibrium of the game, we must have $f_1 = f_2$; otherwise, the exchange with the lower fee benefits from increasing its fee to the fee of the other exchange minus ε (where ε is a sufficiently small positive real number) and still get all of the advertisers. Finally, it is easy to see that $f_1 = f_2 = 0$ is the only equilibrium of the game. If the fees are larger than zero (i.e., $f_1 = f_2 > 0$), at least one exchange benefits from lowering its fee by ε to get all of the advertisers.

Alternative Timelines

Exchanges decide the auction format before the publisher sets the reserve prices. Next, we explore an alternative timeline of the game. More specifically, we focus on step 3 (where the publisher sets the reserve prices) and step 4 (where the exchanges decide the auction formats) of the main model (see the "Timeline" subsection for the main timeline). We show that Proposition 6 (the analog of Proposition 4) is robust under a change in the order of these steps.

Proposition 6: Under header bidding and for any values of $n_1, n_2 > 0$, there is a unique equilibrium where both exchanges use first-price auctions.

The proof of Proposition 6 is available in the Web Appendix.

Advertisers choose an exchange after exchanges decide the auction format. In this subsection, we continue with the timeline of the previous subsection but with an additional change. We assume that the advertisers decide which exchange to join (step 2 in the "Timeline" subsection) after the exchanges decide their auction format and before the publisher sets the reserve prices. We show that Proposition 7 (the analog of Proposition 4) holds in this new timeline.

Proposition 7: Under header bidding, either both exchanges use first-price auctions or all advertisers join the same exchange in equilibrium.

The proof of Proposition 7 is available in the Web Appendix.

Appendix B: Horizontally Differentiated Exchanges

Buyer-Side Fees with Loyals and Switchers

In Proposition 5, we see that under header bidding with first-price auctions, the buyer-side fees decline to zero. If we relax

the assumption of nonnegative fees in the main model and allow the exchanges to set negative fees (i.e., paying the advertisers to join), then in equilibrium the fees become negative down to the point where the exchanges make zero total profit. In practice, these type of equilibria do not usually occur because there is some differentiation between exchanges that prevents the Bertrand-type competition. Next, we extend our model to consider horizontally differentiated exchanges and allow negative buyer-side fees to investigate the fee structure that results from these changes. The goal is to establish the robustness of the main results under this extension.

To model horizontal differentiation, we use a standard loyal/switcher model (see, e.g., Iyer, Soberman, and Villas-Boas 2005; Narasimhan 1988). A loyal advertiser to an exchange is an advertiser that never goes to the other exchange (i.e., it goes either to the exchange it is loyal to or to none of the exchanges [if the expected payoff from joining the exchange it is loyal to is negative]). A switcher is an advertiser that can go to either exchange depending on the expected payoff.

For illustrative purposes for four advertisers, we assume that each exchange has one loyal advertiser and that there are two switchers who can choose between the exchanges. We assume that the buyer-side fees can potentially be negative (i.e., the exchanges can pay the advertisers to join them, but they only do so if the final expected utility is positive). We analyze the waterfalling and the header bidding games subsequently. The proofs of the propositions are available in the Web Appendix.

Waterfalling. Similarly to other works that use the loyal/switchers model, the game does not have a pure-strategy equilibrium. This is because, on the one hand, the exchanges want to lower the fees to attract the switchers, but, on the other hand, they want to increase the fees to extract more surplus from the loyal advertisers. We present a mixed-strategy equilibrium in Proposition 8.

Proposition 8: Under waterfalling, there is a mixed equilibrium where Exchange 1 (the first one in the waterfall) with probability $p = \frac{27(6f+5)}{8(708f+125)}$ sets the buyer-side fee to be $f_1 = \frac{225}{8,192}$, while with probability $1-p$ it chooses a fee drawn from the distribution with CDF

$$G_1(x) = \frac{72(708f + 125)(917f + 8,192x - 75)}{(5,502f + 865)(8,496f + 49,152x + 185)}$$

and domain $\left[\frac{75-917f}{8,192}, \frac{1,315}{49,152}\right)$, where f is the seller-side fee. Exchange 2 with probability $q = \frac{35}{16,506f+2,630}$ sets the buyer-side fee to be $f_2 = \frac{125}{4,096}$ while with probability $1-q$ it chooses a fee drawn from the distribution with CDF

$$G_2(x) = \frac{(8,253f + 1,315)(5,502f + 49,152x - 635)}{(5,502f + 865)(8,253f + 49,152x - 185)}$$

and domain $\left[\frac{635-5,502f}{49,152}, \frac{125}{4,096}\right)$. In equilibrium, exchanges are indifferent between first-price and second-price auctions.

Note that in contrast to Example 1, in the equilibrium described by Proposition 8 both exchanges have a positive number of advertisers for any realization of the fees.

Header bidding. In the following proposition, we describe the mixed equilibrium for the buyer-side fees under header bidding. In contrast to the waterfalling case, here we have a symmetric mixed equilibrium.

Proposition 9: Under header bidding, there is a mixed equilibrium where both exchanges choose a fee drawn from the distribution with CDF $G(x) = \frac{98f+960x-13}{98f+640x}$ and domain $\left[\frac{13-98f}{960}, \frac{13}{320}\right]$. In equilibrium, both exchanges use first-price auctions.

Propositions 8 and 9 show the robustness of our results when exchanges are horizontally differentiated and the fees are allowed to be negative. We see that, in equilibrium, while the exchanges are indifferent between first-price and second-price auctions under waterfalling, they uniquely use first-price auctions under header bidding. Proposition 9 also shows that, when the exchanges are horizontally differentiated, their buyer-side revenue does not decline all the way to zero under header bidding.

Multiple segments of advertisers. In the previous subsection, we consider a loyal/switcher model of horizontally differentiated exchanges and show that under header bidding both exchanges will end up with a positive number of advertisers and will use first-price auctions in equilibrium (in contrast to Proposition 7, where there is an equilibrium where all advertisers join the same exchange and the auction format does not matter).

In this subsection, we consider a more general model of horizontally differentiated exchanges and show that even under very mild assumptions of horizontal differentiation, the monopolistic equilibrium where all advertisers join the same exchange is eliminated. As a result, exchanges uniquely use first-price auctions in equilibrium.

In practice, advertisers can have an intrinsic preference for one exchange,[19] but they can also go to another exchange if the benefit is sufficiently high. To capture this intuition, we assume that there are three segments of advertisers: Segment 1, Segment 2, and Segment 0. Segment 1 advertisers have a slight preference $\delta > 0$ for joining Exchange 1 over Exchange 2. For example, a Segment 1 advertiser will choose Exchange 1 as long as its payoff by moving to Exchange 2 will not increase by more than δ. Similarly, Segment 2 advertisers have a slight preference $\delta > 0$ for joining Exchange 2 over Exchange 1. Finally, Segment 0 advertisers do not have any intrinsic preference for an exchange and, all else being equal, they are indifferent between the two exchanges.

In the next subsection, we consider this setting under the timeline of the main model and the timeline described in

[19] For example, an exchange may offer specialized analytics tools that are more valuable to some advertisers than to others. In addition, some advertisers might have already made infrastructure investments with an exchange and the cost of switching to a different exchange may be nonnegligible.

Appendix A, "Exchanges Decide the Auction Format Before the Publisher Sets the Reserve Prices." In the Web Appendix, we consider it under the timeline of Appendix A, "Advertisers Choose an Exchange After Exchanges Decide the Auction Format."

Advertisers choose exchanges before the auction format decision. In this subsection, we consider the two timelines of the "Timeline" subsection in the main text and the "Exchanges Decide the Auction Format Before the Publisher Sets the Reserve Prices" subsection in Appendix A. We show that under both timelines, as long as Segments 1 and 2 are nonempty, the following proposition holds.

Proposition 10: For any $\delta > 0$, both exchanges have a positive number of advertisers and use first-price auctions in equilibrium.

Proposition 10 shows that as long as each exchange is preferred by at least one advertiser, even if the preference is infinitesimal, we get a unique equilibrium in which both exchanges use first-price auctions. In other words, even a mild assumption of horizontal differentiation eliminates the equilibrium where all advertisers join the same exchange and the exchange is indifferent between the auction formats. The proof of Proposition 10 is available in the Web Appendix.

Acknowledgments

The authors thank seminar participants at Duke University, UC Berkeley, London Business School, and Amazon.com for their comments.

Author Contributions

Stylianos Despotakis, R. Ravi, and Amin Sayedi contributed equally to the article.

Associate Editor

Anthony Dukes

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) received no financial support for the research, authorship, and/or publication of this article.

References

Akbarpour, Mohammad and Shengwu Li (2020), "Credible Auctions: A Trilemma," *Econometrica*, 88 (2), 425–67.

Amaldoss, Wilfred, Preyas S. Desai, and Woochoel Shin (2015), "Keyword Search Advertising and First-Page Bid Estimates: A Strategic Analysis," *Management Science*, 61 (3), 507–19.

Amaldoss, Wilfred, Kinshuk Jerath, and Amin Sayedi (2015), "Keyword Management Costs and 'Broad Match' in Sponsored Search Advertising," *Marketing Science*, 35 (2), 259–74.

Berman, Ron (2018), "Beyond the Last Touch: Attribution in Online Advertising," *Marketing Science*, 37 (5), 771–92.

Berman, Ron and Zsolt Katona (2013), "The Role of Search Engine Optimization in Search Marketing," *Marketing Science*, 32 (4), 644–51.

Bigler, Jason (2019), "Rolling Out Its First Price Auctions to Google Ad Manager Partners," Google Ad Manager (September 5), https://www.blog.google/products/admanager/rolling-out-first-price-auctions-google-ad-manager-partners/.

Bruce, Norris I., B.P.S. Murthi, and Ram C. Rao (2017), "A Dynamic Model for Digital Advertising: The Effects of Creative Format, Message Content, and Targeting on Engagement," *Journal of Marketing Research*, 54 (2), 202–18.

Cao, Xinyu and T. Tony Ke (2019), "Cooperative Search Advertising," *Marketing Science*, 38 (1), 44–67.

Choi, Hana and Carl F. Mela (2018), "Display Advertising Pricing in Exchange Markets," working paper, Duke University.

Choi, Hana, Carl F. Mela, Santiago R. Balseiro, and Adam Leary (2020), "Online Display Advertising Markets: A Literature Review and Future Directions," *Information Systems Research*, 31 (2), 556–75.

Choi, W. Jason and Amin Sayedi (2019), "Learning in Online Advertising," *Marketing Science*, 38 (4), 584–608.

Cox, Sam (2019), "Simplifying Programmatic: First Price Auctions for Google Ad Manager," Google Ad Manager (March 6), https://www.blog.google/products/admanager/simplifying-programmatic-first-price-auctions-google-ad-manager/.

Desai, Preyas S., Woochoel Shin, and Richard Staelin (2014), "The Company That You Keep: When to Buy a Competitor's Keyword," *Marketing Science*, 33 (4), 485–508.

Despotakis, Stylianos, R. Ravi, and Kannan Srinivasan (2021), "The Beneficial Effects of Ad Blockers," *Management Science*, 67 (4), 2096–125.

Dukes, Anthony J., Qihong Liu, and Jie Shuai (2019), "Interactive Advertising: The Case of Skippable Ads," https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3169629.

Gritckevich, Aleksandr, Zsolt Katona, and Miklos Sarvary (2018), "Ad Blocking," Columbia Business School Research Paper No. 18–11.

Gubbins, Paul (2017), "Things You Should Know About SSP/Exchange Auctions, Today," LinkedIn (October 2), https://www.linkedin.com/pulse/things-you-should-know-sspexchange-auctions-today-paul-gubbins/.

Hoban, Paul R. and Randolph E. Bucklin (2015), "Effects of Internet Display Advertising in the Purchase Funnel: Model-based Insights from a Randomized Field Experiment," *Journal of Marketing Research*, 52 (3), 375–93.

Hu, Yu (Jeffrey), Jiwoong Shin, and Zhulei Tang (2015), "Incentive Problems in Performance-based Online Advertising Pricing: Cost Per Click vs. Cost per Action," *Management Science*, 62 (7), 2022–38.

Insider Intelligence (2016), "Header Bidding Gains Momentum, Drives Up Publisher and Revenue," *Business Insider* (May 3), https://www.businessinsider.com/header-bidding-gains-momentum-drives-up-publisher-ad-revenue-2016-5.

Iyer, Ganesh, David Soberman, and J. Miguel Villas-Boas (2005), "The Targeting of Advertising," *Marketing Science*, 24 (3), 461–76.

Jerath, Kinshuk, T.T. Ke, and Fei Long (2018), "The Logic and Management of 'Digital Co-op' in Search Advertising," working paper, Columbia University.

Jerath, Kinshuk, Liye Ma, Young-Hoon Park, and Kannan Srinivasan (2011), "A 'Position Paradox' in Sponsored Search Auctions," *Marketing Science*, 30 (4), 612–27.

Johnson, Garrett (2013), "The Impact of Privacy Policy on the Auction Market for Online Display Advertising," Simon School Working Paper No. FR 13-26.

Katona, Zsolt and Miklos Sarvary (2010), "The Race for Sponsored Links: Bidding Patterns for Search Advertising," *Marketing Science*, 29 (2), 199–215.

Katona, Zsolt and Yi Zhu (2017), "Quality Score That Makes You Invest," working paper.

Kershaw, Tom (2017), "Greater Transparency & Choice in Auction Dynamics," Rubicon Project (December 11), https://rubiconproject.com/insights/thought-leadership/greater-transparency-choice-in-auction-dynamics/.

Krishna, Vijay (2009), *Auction Theory*. Cambridge, MA: Academic Press.

Kuksov, Dmitri, Ashutosh Prasad, and Mohammad Zia (2017), "In-Store Advertising by Competitors," *Marketing Science*, 36 (3), 402–25.

Lambrecht, Anja and Catherine Tucker (2013), "When Does Retargeting Work? Information Specificity in Online Advertising," *Journal of Marketing Research*, 50 (5), 561–76.

Long, Fei, Kinshuk Jerath, and Miklos Sarvary (2018), "The Informational Role of Sponsored Advertising on Online Retail Marketplaces," working paper, Columbia University.

Lu, Shijie, Yi Zhu, and Anthony Dukes (2015), "Position Auctions with Budget Constraints: Implications for Advertisers and Publishers," *Marketing Science*, 34 (6), 897–905.

Milgrom, Paul R. and Robert J. Weber (1982), "A Theory of Auctions and Competitive Bidding," *Econometrica*, 50 (5), 1089–122.

Myerson, Roger B. (1981), "Optimal Auction Design," *Mathematics of Operations Research*, 6 (1), 58–73.

Narasimhan, Chakravarthi (1988), "Competitive Promotional Strategies," *Journal of Business*, 427–49.

OpenX (2019), "Identifying First-Price Auctions," (January 16), https://docs.openx.jp/demandpartners/first-price-auctions.html.

PubMatic (2021), "Support for First-Price Auctions FAQs," (accessed July 15, 2021), https://community.pubmatic.com/display/SSP/Support+for+First-Price+Auctions+FAQs.

Rafieian, Omid (2019), "Optimizing User Engagement Through Adaptive Ad Sequencing," Technical Report, University of Washington.

Rafieian, Omid (2020), "Revenue-Optimal Dynamic Auctions for Adaptive Ad Sequencing," Technical Report, Cornell University.

Rafieian, Omid and Hema Yoganarasimhan (2021), "Targeting and Privacy in Mobile Advertising," *Marketing Science*, 40 (2), 193–218.

Sayedi, Amin (2018), "Real-Time Bidding in Online Display Advertising," *Marketing Science*, 37 (4), 553–68.

Sayedi, Amin, Kinshuk Jerath, and Marjan Baghaie (2018), "Exclusive Placement in Online Advertising," *Marketing Science*, 37 (6), 970–86.

Sayedi, Amin, Kinshuk Jerath, and Kannan Srinivasan (2014), "Competitive Poaching in Sponsored Search Advertising and Its Strategic Impact on Traditional Advertising," *Marketing Science*, 33 (4), 586–608.

Shin, Woochoel (2015), "Keyword Search Advertising and Limited Budgets," *Marketing Science*, 34 (6), 882–96.

Shterev, Boris (2020), "First-Price Auction—What Does It Mean for Publishers?" PubGalaxy (April 24), https://www.pubgalaxy.com/blog/ad-tech/first-price-auction-what-does-it-mean-for-publishers/.

Sluis, Sarah (2017a), "Big Changes Coming to Auctions, as Exchanges Roll the Dice on First-Price," Ad Exchanger (September 5), https://adexchanger.com/platforms/big-changes-coming-auctions-exchanges-roll-dice-first-price/.

Sluis, Sarah (2017b), "More on the Widespread Practice Behind the Guardian's Lawsuit vs. Rubicon Project," Ad Exchanger (March 30), https://adexchanger.com/ad-exchange-news/explainer-widespread-fee-practice-behind-guardians-lawsuit-vs-rubicon-project.

Sluis, Sarah (2017c), "Rubicon Project Eliminates Buy-Side Fees," Ad Exchanger (November 2), https://adexchanger.com/platforms/rubicon-project-eliminates-buy-side-fees.

Sluis, Sarah (2019), "Google Switches to First-Price Auction," Ad Exchanger (March 6), https://adexchanger.com/online-advertising/google-switches-to-first-price-auction/.

Tunuguntla, Srinivas and Paul R. Hoban (2021), "A Near-Optimal Bidding Strategy for Real-Time Display Advertising Auctions," *Journal of Marketing Research*, 58 (1), 1–21.

Wilbur, Kenneth C. and Yi Zhu (2009), "Click Fraud," *Marketing Science*, 28 (2), 293–308.

Zawadziński, Maciej (2018), "Waterfalling, Header Bidding and New Auction Dynamics," Clearcode (January 31), https://clearcode.cc/blog/sequential-auctions-header-bidding-first-price-second-price-auctions/.

Zeithammer, Robert (2019), "Soft Floors in Auctions," *Management Science*, 65 (9), 4204–21.

Zhu, Yi and Kenneth C. Wilbur (2011), "Hybrid Advertising Auctions," *Marketing Science*, 30 (2), 249–73.

Zia, Mohammad and Ram C. Rao (2019), "Search Advertising: Budget Allocation Across Search Engines," *Marketing Science*, 38 (6), 1023–37.